DOCUMENT RESUME

ED 410 864 HE 030 460

AUTHOR Tamada, Mike; Inman, Claudia
TITLE Survival Analysis of Faculty Retention Data: How Long Do They Stay? AIR 1997 Annual Forum Paper.
PUB DATE 1997-05-00
NOTE 36p.; Paper presented at the Annual Forum of the Association for Institutional Research (37th, Orlando, FL, May 18-21, 1997).
PUB TYPE Numerical/Quantitative Data (110) -- Reports - Research (143) -- Speeches/Meeting Papers (150)
EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *College Faculty; Decision Making; *Faculty College Relationship; *Faculty Mobility; Higher Education; *Sex Differences; Statistical Analysis; Tables (Data); Teacher Attitudes; *Teacher Persistence; Tenure
IDENTIFIERS *AIR Forum; *Survival Analysis

ABSTRACT

This study applied survival analysis methodology to faculty retention data in order to examine ways to measure faculty retention and determine whether men and women have different "survival times." The study at a selective, private liberal arts college first used college catalogs to identify 339 full-time tenure-track faculty who had begun working in 1960 through 1994 and to determine their final year at the college. Variables included Ph.D. or "All But Dissertation" (A.B.D.) status when hired, year of Ph.D., entry rank, year of full-time tenure-track status, and tenure status upon entry. Application of the survival analysis techniques indicated that faculty who arrived in the earlier years (1960s and 1970s) had much lower durations than faculty who arrived later. Comparison of male and female retention rates indicated that female retention rates were essentially the same as male rates. Details of the methodology and its limitations are discussed. Twelve tables of data are appended. (Contains 14 references.) (JLS)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED 410 864

SURVIVAL ANALYSIS OF FACULTY RETENTION DATA:

HOW LONG DO THEY STAY?

Mike Tamada and Claudia Inman
Occidental College
Los Angeles, CA

Presented at the Association for Institutional Research
37th Annual Forum, Orlando FL
May, 1997

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)
✓ This document has been reproduced as received from the person or organization originating it.
☐ Minor changes have been made to improve reproduction quality.
• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

AIR

ABSTRACT:

This is an introduction to survival analysis, applied to faculty retention data. If a college is concerned about how long its faculty stay, especially women faculty (perhaps for Title IX purposes), then two questions arise: how to measure retention, and how to discern whether men and women have different "survival times". A set of special statistical techniques known as "survival analysis" is useful for answering questions such as these. We informally describe the techniques used and why they are useful. Looking at all tenure track faculty from 1960, we found that women's retention was essentially the same as men's; however the data do not tell us the reason for departures.

IR 030 460

ACKNOWLEDGMENTS:

We thank Amber Reisz for invaluable research assistance, and the California Association for Institutional Research Conference Presentations Committee for helpful comments. This work was supported in part by a summer research fellowship from the National Science Foundation and California Alliance for Minority Participation (NSF-CAMP).

BEST COPY AVAILABLE

for Management Research, Policy Analysis, and Planning

This paper was presented at the Thirty-Seventh Annual Forum of the Association for Institutional Research held in Orlando, Florida, May 18-21, 1997. This paper was reviewed by the AIR Forum Publications Committee and was judged to be of high quality and of interest to others concerned with the research of higher education. It has therefore been selected to be included in the ERIC Collection of Forum Papers.

Jean Endo
Editor
AIR Forum Publications

SECTION I. INTRODUCTION

How long do faculty stay at a college, and do male and female faculty have different "survival times"? Questions such as these involve data with special characteristic, and special statistical methods, known variously as "survival analysis" or "duration analysis" or "analysis of failure time data" are required. This is an informal introduction to survival analysis, applied to faculty retention data.

The use of survival analysis is gradually becoming more widespread in the social sciences. See, for example, a review of the econometric literature by Kiefer (1988), econometric work by Heckman and Borjas (1980), and introductory articles in the psychology literature by Morita et al (1989) and Singer and Willet (1991). Institutional researchers are starting to use survival analysis; a prominent example is the recent article for the AIR Professional File by Ronco (1996) describing the "competing risks" model. Also, at the 1995 California Association for Institutional Research Conference, Garcia (1995) used life table methodology to track student retention and graduation rates. Statistical packages such as SPSS are gradually being given more and more powerful survival analysis capabilities, enabling researchers to more easily carry out such analyses.

Survival analysis is useful for answering questions involving some sort of duration; the question could be the survival time of cancer patients, the duration of unemployment spells, the age at which people first get married, the retention of faculty -- in short, any question involving the length of time that passes until a certain event occurs (death, employment, marriage, termination or exit from the school, and so on).

But most undergraduate and even graduate statistics courses in most disciplines do not cover survival analysis. This paper will introduce a few of the concepts of survival analysis, starting with the basic definitions and moving up to regression analysis of survival data. These techniques will be applied to faculty retention data, in particular to test the null hypothesis that male and female faculty at a private liberal arts college have equal survival times.

PRIOR RESEARCH ON FACULTY RETENTION

Although many studies of gender differences among faculty exist (see Dwyer et al (1991) for a survey), longitudinal studies of faculty retention are much rarer. Most earlier studies seem to have found higher mobility rates (i.e. lower retention rates) for female faculty than for male faculty. However, these either were comparative rather than longitudinal, or dealt with only a specific subset of faculty, such as psychology faculty (Rosenfeld and Jones 1986) or part-time faculty (Tuckman and Tuckman 1981). This study, though covering only one school, covers all tenure-track faculty in all fields and follows them longitudinally. By covering the faculty at one school only this study does lose generality, but at the same time avoids the complications involved in comparing faculty at research institutions with those at teaching institutions, and comparing faculty of widely divergent backgrounds and quality levels. Moreover, this study illustrates how a wider ranging study could be performed, if longitudinal data on a variety of institutions were gathered.

Ashenfelter and Card (1996) are working with TIAA/CREF and the Princeton Retirement Survey to create a database with which they can study faculty retirement using

survival analysis techniques. However this database's usefulness in studying the retention of junior faculty will be somewhat limited because it will not include professors who left their schools prior to 1986.

In Section II we will introduce some of the fundamental concepts used in survival analysis: survivor functions, hazard functions, and censoring. In Section III we will describe the issue being researched, namely faculty retention by gender, and describe the data set. In section IV we will describe simple techniques for analyzing the data, such as using life tables to look at the survivor functions and performing log-rank tests for differences between the genders. In section V we will discuss more complex techniques, such as Cox's proportional hazards regression model.

SECTION II. FUNDAMENTAL CONCEPTS

LIFE TABLES AND SURVIVOR FUNCTIONS

Some of the most fundamental concepts of survival analysis can be illustrated with a life table, similar to the ones used by actuaries and demographers. Suppose that in the year 1900, 100 children were born in Costa Mesa, California. By 1901, let us suppose that 90 of them were still alive; by 1902, 80 of them were still alive; and by 1903, 70 survived. We could begin constructing a life table that would look like the one in Exhibit 1.

The first five columns are fairly self-explanatory. The "observed survivor function" simply tells us, for any given year, what percent of the population is still surviving. This

function will always be non-increasing, as long as we are dealing with standard single-event survival models.

HAZARD FUNCTIONS

However, researchers will often choose not to focus on the survivor function, but instead will focus on the "hazard function" -- the percent of REMAINING SURVIVORS (not the percent of the total) who die in a given year. Notice that in Exhibit 1, even though a constant number of ten people are dying each year, and thus a constant 10% of the population is dying each year, the hazard rate is INCREASING. In the second year, the ten deaths represent one-NINTH of the survivors, and in the third year one-eighth (1).

When dealing with human mortality, the "mortality rate" is simply another name for the "hazard rate," i.e. the value of the hazard function.

Mathematically, the hazard function can be derived from the survivor function and vice-versa (2). But in practical terms for researchers, hazard functions are often more convenient to study. One reason is that survivor functions, when graphed, all look pretty much alike -- they are all downward-sloping. It is often difficult to distinguish between different survivor functions graphically, and to deduce what the graph is telling us.

Hazard functions in contrast will typically have very different appearances from population to population, or from model to model. The researcher can more easily interpret and tell a story about a given hazard function. For example, if I asked you what the hazard function for human beings looked like (realistically, not using the fake data in Exhibit 1), after a little thought you would probably realize that it is U-shaped: the mortality rate for

infants is relatively high, then it falls for children and young adults, then it rises continually for older people.

In contrast to the hazard function, it would be difficult for you to tell me much about human beings' survivor function, except that it is downward-sloping.

Radioactive decay provides another example of a hazard function. How many Cesium 137 atoms are left after a period of time? Radioactive decay is usually assumed to be constant, and since Cesium 137 has a half-life of 30.0 years, about 2.3% of the atoms will decay per year. This is a CONSTANT hazard rate, with 2.3% decay per year.

Hazard functions often exhibit "negative time dependence," that is, the hazard rate decreases over time. Unemployment spells often are an example: quite a few unemployment spells end after two or three months, but by the time an unemployment spell has lasted, say sixty months, the job-seeker's probability of finding a job in the next month is quite small -- i.e. his or her hazard rate is low. We will see that the hazard function for faculty typically exhibits positive time dependence initially, but after a few years the exhibits negative time dependence. After a professor has been around for 15 years, the probability that she will leave in the 16th year is low.

CENSORING

Survival data frequently are "censored," meaning that the true value of an subject's survival time is unknown, except that it exceeds a certain value. Here is an example of censoring: as of 1996, any professor who arrived in 1994 would have completed two years and the value of their duration variable would be 2. But these faculty are very different from

faculty who arrived in, say, 1968 and left in 1970. Both have durations of 2 years, but the 1994 cohort of faculty will ultimately have durations GREATER than 2 -- but we do not know what their final, true, duration will be. Thus returning faculty are all censored (3). We only know the true durations of faculty who have arrived AND left the school.

How do we deal with censoring? Clearly, it is undesirable to take the duration variables at face value and to consider the 1994 faculty to have durations of 2 years. One possibility is to drop the censored subjects from the data set. This usually creates two major problems however. First, the data set may as a result shrink to an unacceptably small number of subjects. Second, the sample will probably be biased -- because professors with long durations are especially likely to be censored, and these long-lived faculty are thus getting dropped from the sample. The sample will be biased towards short-lived faculty.

A better way of dealing with censoring is to use the survival analysis techniques which have been developed over the years to deal with the problem of censoring. These will be explained after we describe the data set.

SECTION III. EXAMINING FACULTY RETENTION

THE ISSUES

This study covers faculty at a selective private liberal arts college. During the late 1980s and early 1990s there seemed to be an unusually large number of junior female faculty who left the school, for various reasons. Also, in the early 1990s the school appointed a professor to be its Title IX coordinator. Thus questions of faculty retention,

especially female faculty retention, arose. Although the school has had a good record of hiring female faculty, and although the proportion of women in the faculty has been rising, there was still the question of whether these newly-hired women faculty were actually staying at the school.

Unlike the situation with students, whose "survival" can be measured with indicators such as graduation and retention rates, there are no widely used overall measures of faculty retention, with the exception of tenure and promotions. However, faculty in general have to stay about 6 years before they can get tenure -- and thus information on tenure will not cover faculty in the first three or four years at the school. Many professors were simply too new to be eligible for tenure, others left before they became eligible. Also, tenure doesn't tell us how long the professor actually stayed at the school; it merely tells us that they stayed long enough to get tenure.

A better measure of retention is to literally count how many years each professor stayed at the school.

DATA

We used the college's catalogs to identify 339 full-time tenure-track faculty who had started working in 1960 up to 1994, and to determine their final year at the college. Many of them of course are still at the college.

The catalog also supplied us with the following variables: PhD/ABD status when hired, year of Phd, department, entry rank, year of full-time tenure-track status (some professors started as adjunct or visiting faculty), and tenure status upon entry (a few

professors enter with tenure in hand). We also collected data on years to tenure and to full professorship, but those variables are not used in this study.

The catalog did not directly supply us with gender information, but by looking at the names and consulting with veteran employees we were able to determine the gender of all but two faculty. We do not have ethnicity information, especially for faculty from the 1960s.

We do not have information on the REASON for exit; the professor may have left due to a better offer elsewhere, or may have been turned down for tenure or contract renewal. Thus this study only measures overall retention; it does not measure retention of "desirable" faculty, or the rate at which "undesirable" faculty were gotten rid of.

Descriptive statistics for the data set are in Exhibit 2.

If there were no censoring problems, we could simply find the mean duration of male and female professors, do a t-test and be done. We could also do linear regressions to see if other variables affect duration.

However, our data our heavily censored -- about 140 of the 339 professors in our sample are still at the school and thus we do not know their ultimate duration. Thus we utilized survival analysis.

SECTION IV. SIMPLE COMPARISONS

SURVIVOR FUNCTIONS

Initial analyses of our data quickly revealed that faculty who arrived in the earlier years -- the 1960s and 1970s -- had much lower durations that faculty who arrived later. (We later

performed a log-rank test showing a large and highly significant difference.) Thus we decided to split the data set, since it seemed apparent that the faculty who arrived in 1980-94 had survival and hazard rates which were different from the 1960-79 faculty. Also, since most of the earlier faculty were men, a simple comparison of men vs. women would tend to show men having a low retention rate simply due to the fact that so many of them arrived during the years when retention rates were low.

Exhibit 3 shows life tables for the faculty who arrived from 1960 to 1979, and the faculty who arrived from 1980 to 1994. The estimated survivor functions are calculated using Kaplan-Meier (also known as product limit) estimates. Notice that the censored observations are utilized "for as long as they can" -- that is, if a professor has been at the school for three years and is still there right now, we do not know her ultimate duration. But we do know that she did not "attrit" (that is, leave the school) after her first or second years, and so she does contribute to the calculation of the one- and two-year retention rates. Most full-featured statistical packages, such as SPSS, will calculate life tables and survivor and hazard functions.

Exhibit 4 graphs the observed survivor rates and Exhibit 5 the observed hazard rates for the 1960-79 and 1980-94 faculty.

We do not know the explanation for the very high attrition rates of the 1960-79 faculty. A non-trivial proportion (one out of nine) only lasted one year. One potential factor is that the school offered only 1-year contracts to new faculty for much of that period. However, it seems unlikely that this is a complete explanation: the vast majority of faculty hired in recent years would probably stay longer than a year even if they were limited to 1-year

contracts. Possibly faculty quality became higher in the 1980s and 1990s, and new faculty are more likely to qualify for contract renewal, tenure, etc.

The hazard functions in Exhibit 5 of course show the very high hazard rates that the early professors experienced, especially in their first few years. Post-1979 professors in contrast have a very low hazard rate their first two years -- 94% of stayed for at least their third year -- and even when their hazard rate increases it still lower than that of the pre-1979 professors.

In addition it is interesting to note that the hazard rates for the pre-1979 professors peaked in their 4th and 7th years -- not too surprising given the timing of tenure decisions and contract renewals. The hazard for the post-1979 professors peaks in their 5th and 8th years, which is somewhat surprising. Possibly more tenure decisions are getting deferred or delayed in recent years. For all professors, the 8th year seems to be the cutoff point -- if a professor has stayed for 8 years, the chances are quite good that he or she will be back for the 9th and subsequent years. This same phenomenon can be observed in the survival graphs in Exhibit 4 -- after the 8th year the survival curves flatten out.

Exhibits 6 and 7 show the life tables for male and female faculty who entered after 1979. Exhibit 8 shows a graph of their observed survivor rates. It appears that men have a slightly higher survivor or retention rate than women, but it partly depends on how where one measures the survivor rate -- for example, women only have an 84% four-year retention rate whereas men have a 91% retention rate. But women have a 58% thirteen-year retention rate, close to the 59% men's rate. On the whole the differences do not seem terribly large -- but how can we tell what "large" is? To some degree this is a decision for policy-makers to

decide. But we can also make an overall comparison of the two survival functions, and measure the statistical significance of the difference. A simple way of testing for the difference between two survival functions is to perform a log-rank test.

LOG RANK TESTS

Log-rank tests are relatively simple to perform (statistical packages such as SPSS will perform these tests). They can be interpreted as a generalization of rank tests such as the Wilcoxon test; essentially the number of attritions in a given period is compared to the number of attritions expected under the null hypotheses. See, for example, Kalbfleisch and Prentice (1980) for a discussion and derivation.

The log-rank test yields a statistic which is distributed as chi-squared, with r-1 degrees of freedom, where r is the number of samples being compared. In our case, we have two post-1979 samples, men and women. The log-rank test was significant at the p=36% level, nowhere close to the standard significance levels and suggesting that the differences between men's and women's survival times could have been caused by random variation.

For the 1960-79 samples, men actually seemed to have lower survival rates than women. However a log-rank test performed on these samples again showed no significant differences.

A log-rank test comparing all post-1979 faculty to all 1960-79 faculty was highly significant however, with p well below 1/10 of 1%.

The log-rank test has an important weakness in that it simply compares two (or more) entire samples. It does not take into account the effects of other variables, such as PhD/ABD

status, entering rank, entering tenure status, and time trends. To control for these other variables, a multivariate approach is preferable.

SECTION V. MORE COMPLEX TESTS

THE COX PROPORTIONAL HAZARDS MODEL

There are several different regression models that can be applied to survival data. Many of them are based on parametric hazard functions; that is, one has to assume that the population has an underlying hazard function with a specific functional form. The simplest such functional form would be the exponential model -- in the exponential model, the hazard rate is constant (that is, a constant proportion, h, of the population exits each period, and the surviving population thus declines exponentially) and a regression might estimate the value of h, as well as the value of the slope parameters of the righthand side variables used in the regression.

Few survival processes have such simple functional forms -- typically the hazard rate will vary with the subject's duration. For such situations there are many more complex parametric models which can be used. Some of them can flexibly fit data with positive time dependence, negative time dependence, or both.

In our case however, we were unwilling to make prior assumptions about the shape of the hazard function, and thus unwilling to choose one specific parametric model.

The "Cox proportional hazards regression model" is a regression model frequently used in such situations. It does not make prior assumptions about the shape of the hazard

function -- the "baseline" hazard function is estimated from the data. It does however assume that all the right hand side variables affect the hazard function proportionately. For example, a change in the value of one righthand side variable might double the entire hazard function; a change in another variable might reduce the entire hazard function. The impact of a righthand side variable is assumed to always be a proportional change in the entire hazard function (4).

Some packages such as the Windows version of SPSS can perform Cox proportional hazards regressions, as can many econometric packages. The coefficients cannot be calculated directly; iterative maximum likelihood techniques are necessary, just as with logit (also known as logistic) regressions.

We ran the regression with several different sets of variables; in all of them gender had only a very small coefficient and was nowhere close to significance at the 5% or even 10% level. We did find however that faculty who entered with a PhD had significantly higher survival rates than faculty who entered ABD and faculty who entered with tenure also had higher survival rates (not surprisingly). There was some evidence that faculty who started as adjuncts also had higher survival rates (however remember that this sample is of tenure track faculty only, and only a small proportion of adjuncts are able to switch into a tenure track position). And of course the post-1979 faculty had much higher survival rates. The professors' departments did not seem to affect survival rates. The results from an illustrative regression are in Exhibit 9. Remember that the dependent variable is the hazard rate, so the negative coefficient on post-1979 faculty means that they have LOWER hazard rates, and thus HIGHER retention and survival rates.

SOME REGRESSION DIAGNOSTICS

There are alternative ways of running the regression, for example the sample can be split into subsamples called strata. Each stratum has its own baseline hazard function, which as before is completely flexible, without any parametric assumptions. However all strata share the same righthand side variables and the same slope coefficient.

How should we decide whether we need to split the sample into strata? More generally, what sorts of regression diagnostics are available, so that we can evaluate the "goodness of fit" of the regression?

First, the bad news. There is no equivalent to the R^2 or the mean squared error that can be used to evaluate Ordinary Least Squared regressions. One can perform a log-likelihood test (which is distributed as a chi-squared statistic) which compares the overall fitted regression to the null regression -- but as with OLS regressions, almost any sort of reasonable righthand side variables will give extremely significant results, and thus one doesn't get a strong sense of how well the regression fit the data. Some pseudo-R^2 formulas based on the change in the log-likelihood have been suggested.

The good news: there are several graphical techniques for evaluating the results of survival regressions. However they are very heuristic in nature; there do not seem to be any fixed formulas for defining when a fit is "good" or "bad"; rather one simply looks at the graph and tries to decide if the fit is good enough. Also, most statistical packages will not produce these graphs for you; you have to download the parameters and data and produce the graphs yourself.

Here is a brief description of a couple examples of these graphical regression diagnostic techniques. One standard technique is the "log-minus-log" plot: a plot of the logarithm of minus the logarithm of the estimated survival functions of the possible strata, plotted with duration on the horizontal axis. In other words, plot ln(-ln(S(t)) against t, where S(t) is the estimated survival rate at time t. (Remember that survival rates are always between 0 and 1; thus the logarithm of the survival function will always be negative. The log-minus-log plot uses the logarithm of MINUS this logarithm.)

When the different strata are plotted on the log-minus-log plot, their plotted curves should ideally stay roughly the same distance from each other. If they do not have this constant separation, then the proportional hazards assumption may be violated, and the regression should be stratified (rather than using the stratum variable as a righthand side variable). Exhibit 10 shows an example of a log-minus-log plot, with the sample stratified by pre-1979 (actually 1960-79) and post-1979 status.

The two curves show a certain amount of change in their distance from each other, and they even cross at year 7. There may not be any exact guidelines for deciding when to stratify, but this would seem to be a situation where stratification is called for. (The stratified regressions gave results very similar to the ones in Exhibit 9.)

Another diagnostic device is the "generalized residual," a concept suggested by Cox and Snell (1968). In the context of survival analysis, generalized residuals are generated by calculating the "integrated hazard" -- the sum, across time, of the values of the hazard function (or the integral with respect to time if continuous time is being used). As Kiefer (1988) notes, "the integrated hazard does not have a particularly convenient interpretation,"

but "it is the basic ingredient in a variety of specification checks." For the Cox proportional hazards model, a generalized residual for a duration t can be calculated by taking the integrated hazard at time t and multiplying it by the exponent of the product of righthand side variables and their coefficients (i.e. e(t) = H(t)exp(xb) where e(t) is the generalized residual for time t, H(t) is the integrated hazard at time t, xb is the vector product of the righthand side variables and their coefficients, and exp() is the exponential function). These residuals can be plotted with a residual of size r on the horizontal axis and the logarithm of the proportion of residuals greater than r on the vertical axis. The resulting plot, if the regression has a good fit, should ideally follow the 45-degree line from the origin. See also Crowley and Hu (1977) for a discussion and example.

Exhibit 11 shows the plot of the generalized residuals from an unstratified regression. Again there seem to be no hard-and-fast formulas for determining when the residuals are sufficiently close to the ideal. However the graph in Exhibit 11 seems to exhibit a good fit.

Exhibit 12 shows the plot of the generalized residuals from the same regression, stratified by pre-1979 and post-1979 status. If anything these residuals seem to have a worse fit that those from the unstratified regression, which seems counterintuitive. Again it is not clear if the generalized residuals in this graph could be considered to be "close enough" to the 45-degree line.

Thus the results from the log-minus-log and generalized residual graphs are not definitive, but do not seem to indicate a gross lack of goodness of fit in the regression.

CONCLUSION

There are many other statistical techniques used in survival analysis, but this paper has provided an introduction. It seems safe to conclude from the survival graphs, log rank tests, and regression results that the survival rates of male and female faculty did not exhibit large differences in a statistical sense. To decide whether the differences are large enough to worry about in a non-statistical sense is a largely subjective judgement, but the life tables and estimated survival functions at least provide numerical measures for comparing the retention of men and women.

One crucial piece of information that our data set does not provide is the reason for faculty attrition. The college may have deliberately made some professors leave, while it may have wished to retain other professors who left the school. And of course the reasons for attrition are typically complex and cannot be captured in a single variable -- some professors may have wanted to stay on the whole but some aspect of the school made the job unattractive; some professors may have been deemed desirable by some members of the college community and undesirable by others. The dataset does not provide even a hint of what the reasons for attrition were; it simply records who stayed and who left, and when.

Thus it is possible that a school could still have a problem with retaining female faculty even if their retention rate equaled that of the male faculty. Possibly the males who left were not deemed desirable by the school while the females were -- or vice-versa.

WHERE TO GO FROM HERE

This paper has only discussed single-spell, single-outcome models. Some events such as unemployment or marriage can happen repeatedly to the same person over time. Also sometimes there are multiple possible events which we wish to measure: a student might stay enrolled until he or she eventually graduates, transfers or drops out -- this is the subject of Ronco's AIR Professional File article (1996), and in a life table context, Garcia's CAIR conference presentation (1995).

For people who wish to perform survival analyses of their own, we have found Singer and Willett's (1991 and 1993) articles to be clearly written and easy to understand. Morita et al provide another good, slightly more technical introduction. For a more mathematical approach, Kiefer's survey article (1988) and Heckman's work (1984, e.g.) represent the econometric approach. For a general statistical approach, Kalbfleisch and Prentice's book (1980) is cited extremely often and provides a good but mathematical introduction. It is getting a little dated now, however.

We used an econometrics package called Limdep (ver 6.0) and SPSS for Windows (ver 6.1.2) to perform these calculations. Many lower cost statistical packages do not have the capability of performing survival analysis. On the other hand, if you have discrete time data, Willett and Singer (1993) describe how some survival analysis can be performed simply by doing a series of logit regressions (also known as logistic regressions), which many statistical packages can perform.

Survival analysis will not replace the t-test and the contingency table in terms of being a "must know" statistical technique. But if you have a data situation where you are

measuring time duration, especially in the presence of censoring, then survival analysis comes in handy indeed.

ENDNOTES:

(1) In calculating the observed hazard function, there are some technicalities associated with the question of whether time is being measured as a continuous or discrete variable. Most statistical packages, including SPSS, will assume that time is continuous, and will make adjustments to the calculated hazard function instead of using the simple calculations in Exhibit 1. In this example, we are measuring time in years. But most people do not literally live exactly 1.00 years or 2.00 years and then drop dead. Instead they may die at any age, such as 1.032 or 2.964. But life tables put people into age categories, such as 0 to 1 years, and 1 to 2 years, and do not record the exact age at death. Still, knowing that , for example, in the first year we started with 100 people and ended with 90 people, we might assume that people died at an even rate throughout the year and assume that during that first year the average size of the surviving population was 95. Thus one possible simple adjustment to the hazard rate would be to calculate it as 10/95 instead of 10/100. With discrete time, such adjustments are not necessary -- for example, faculty duration typically can be measured in integer years.

(2) If we assume that survival time is a random variable, and denote the survivor function as S(t), where S denotes the proportion of the population surviving at time t, then the

cumulative distribution function of survival time is F(t) = 1-S(t). If time is continuous (rather than discrete), then the density function of survival time is f(t) = F'(t). And the hazard function is h(t) = f(t)/S(t). Conversely, the survivor function can be derived from the hazard function: S(t) = exp(-int(h(t))) where "int(h(t))" denotes the integral of h(t) from 0 to t.

(3) This is known as "right censoring," where the subject's date of EXIT is unknown. In other types of research, subjects can be "left censored," with the date of ENTRY unknown. For example, if one wishes to measure the life expectancy of AIDS patients from the date of infection (as opposed to the date of diagnosis), many patients will not know their date of infection and thus they will be left censored. If they are still alive, they are also right censored.

Sometimes it is also useful to distinguish between Type I censoring and Type II censoring. Type I censoring occurs when the experiment or observations must end at a certain time, and certain subjects will not have experienced the exit event (death, departure from school, etc.). Type II censoring occurs when the researcher stops collecting observations after a certain NUMBER of exit events, for example after 30 faculty have left the school.

(4) Mathematically, the Cox proportional hazards model assumes that the hazard rate (h), is a function of time (t) and a vector of righthand side variables (x) multiplied by a vector of slope coefficients (b). That is, $h(t,x) = h_o(t)\exp(xb)$, where $h_o(t)$ is the baseline hazard function (the underlying hazard function which applies to all members of the population),

and exp() is the exponential function. Large positive values of x and b, for example would cause the hazard rate h(t,x) to increase, raising the attrition rate. Large negative values of x and b would cause the hazard rate to become smaller (but still positive -- hazard rates have to always be nonnegative by definition).

BIBLIOGRAPHY

Ashenfelter, Orley and Card, David. "Faculty Retirement in the Post-Mandatory Era: Early Findings from the Princeton Retirement Survey," Princeton Conference on Higher Education, March 1996.

Cox, David R. and Snell, E. J. "A General Definition of Residuals," Journal of the Royal Statistical Society, Vol. B30 (May/Aug. 1968), pp. 248-275.

Crowley, John and Hu, Marie. "Covariance Analysis of Heart Transplant Survival Data," Journal of the American Statistical Association, Vol. 72, No. 357 (March 1977), pp. 27-36.

Dwyer, Mary M.; Flynn, Arlene A.; and Inman, Patricia S. "Differential Progress of Women Faculty: Status 1980-1990," Higher Education: Handbook of Theory and Research, Vol. 12 (1991), pp. 173-222.

Garcia, Philip. "California Colleges and University Enrollment Demand: 1994-2005," CAIR Annual Conference, Sacramento CA, November 9, 1995.

Heckman, James J. and Borjas, George J. "Does Unemployment Cause Future Unemployment? Definitions, Questions, and Answers from a Continuous Time Model of Heterogeneity and State Dependence," Economica, Vol. 47, No. 187 (March 1980), pp. 248-283.

Kalbfleisch, John D. and Prentice, Ross L. The Statistical Analysis of Failure Time Data. New York: John Wiley & Sons, Inc. 1980.

Kiefer, Nicholas. "Economic Duration Data and Hazard Functions," Journal of Economic Literature, Vol. 26, No. 8 (June 1988), pp. 646-679.

Morita, June G.; Lee, Thomas W.; and Mowday, Richard T. "Introducing Survival Analysis to Organizational Researchers: A Selected Application to Turnover Research," Journal of Applied Psychology, Vol. 74, No. 2 (April 1989), pp. 280-292.

Ronco, Sharron L. "How Enrollment Ends: Analyzing the Correlates of Student Graduation, Transfer and Dropout with a Competing Risks Model," AIR Professional File, No. 61, Summer 1996.

Rosenfeld, R. A., and Jones, J. A. "Institutional Mobility Among Academics: The Case of Psychologists," Sociology of Education, Vol. 59 (1986), pp. 212-226.

Singer, Judith D. and Willett, John B. "Modeling the Days of Our Lives: Using Survival Analysis When Designing Longitudinal Studies of Duration and Timing of Events," Psychological Bulletin, Vol. 110, No. 2 (1991), pp. 268-290.

Tuckman, B.H. and Tuckman, H.P. "Women as Part-Time Faculty Members," Higher Education, Vol. 10, No. 2 (1981), pp. 169-179.

Willett, John B. and Singer, Judith D. "Investigating Onset, Relapse, and Recovery: Why You Should, and How You Can, Use Discrete-Time Survival Analysis to Examine Event Occurrence," Journal of Consulting and Clinical Psychology, Vol. 61, No. 6 (1993), pp. 952-965.

EXHIBIT 1. A SIMPLE LIFE TABLE

Year	Number of Survivors	Number Who Died During the Year	"Observed Survivor Function" = Percent Surviving	Percent of Total Who Died During the Year	"Observed Hazard" = Percent of REMAINING SURVIVORS Who Died During the Year
1900	100	--	100%	--	--
1901	90	10	90%	10%	10.0%
1902	80	10	80%	10%	11.1%
1903	70	10	70%	10%	12.5%

EXHIBIT 2: DESCRIPTIVE STATISTICS

Number of observations = 339

Mean values of Dummy (Binary) Variables (0=no, 1=yes)

Female	.31
Censored	.35
Adjunct (when hired)	.06
HavePhD	.64

Mean values of Numerical Variables

Yrfulltm (the year that the prof became fulltime tenure track)	1975.29
Duration (no correction made for censored observations)	7.98

EXHIBIT 3: The Survival and Hazard Rates of Faculty Arriving Between 1960-79

Years Completed	Number of Faculty	Terminal Events	Censored	Propn Surviving	Cumul Prop Surv at End	Hazard Rate
1	222	0	0	1.000	1.000	0.000
2	222	26	0	0.883	0.883	0.117
3	196	33	0	0.832	0.734	0.168
4	163	31	0	0.810	0.595	0.190
5	132	21	0	0.841	0.500	0.159
6	111	15	0	0.865	0.432	0.135
7	96	16	0	0.833	0.360	0.167
8	80	11	0	0.863	0.311	0.138
9	69	4	0	0.942	0.293	0.058
10	65	3	0	0.954	0.279	0.046
11	62	2	0	0.968	0.270	0.032
12	60	0	0	1.000	0.270	0.000
13	60	3	0	0.950	0.257	0.050
14	57	0	0	1.000	0.257	0.000
15	57	0	0	1.000	0.257	0.000

EXHIBIT 3: The Survival and Hazard Rates of Faculty Arriving Between 1980-94

Years Completed	Number of Faculty	Terminal Events	Censored	Propn Surviving	Cumul Prop Surv at End	Hazard Rate
1	117	0	0	1.000	1.000	0.000
2	117	2	0	0.983	0.983	0.017
3	115	5	7	0.957	0.940	0.043
4	103	7	8	0.932	0.876	0.068
5	88	8	6	0.909	0.797	0.091
6	74	5	8	0.932	0.743	0.068
7	61	2	5	0.967	0.718	0.033
8	54	5	7	0.907	0.652	0.093
9	42	1	5	0.976	0.636	0.024
10	36	2	2	0.944	0.601	0.056
11	32	1	6	0.969	0.582	0.031
12	25	0	3	1.000	0.582	0.000
13	22	0	5	1.000	0.582	0.000
14	17	1	3	0.941	0.548	0.059
15	13	0	5	1.000	0.548	0.000

EXHIBIT 4: SURVIVAL RATES OF FACULTY ARRIVING 1960-79 and 1980-94

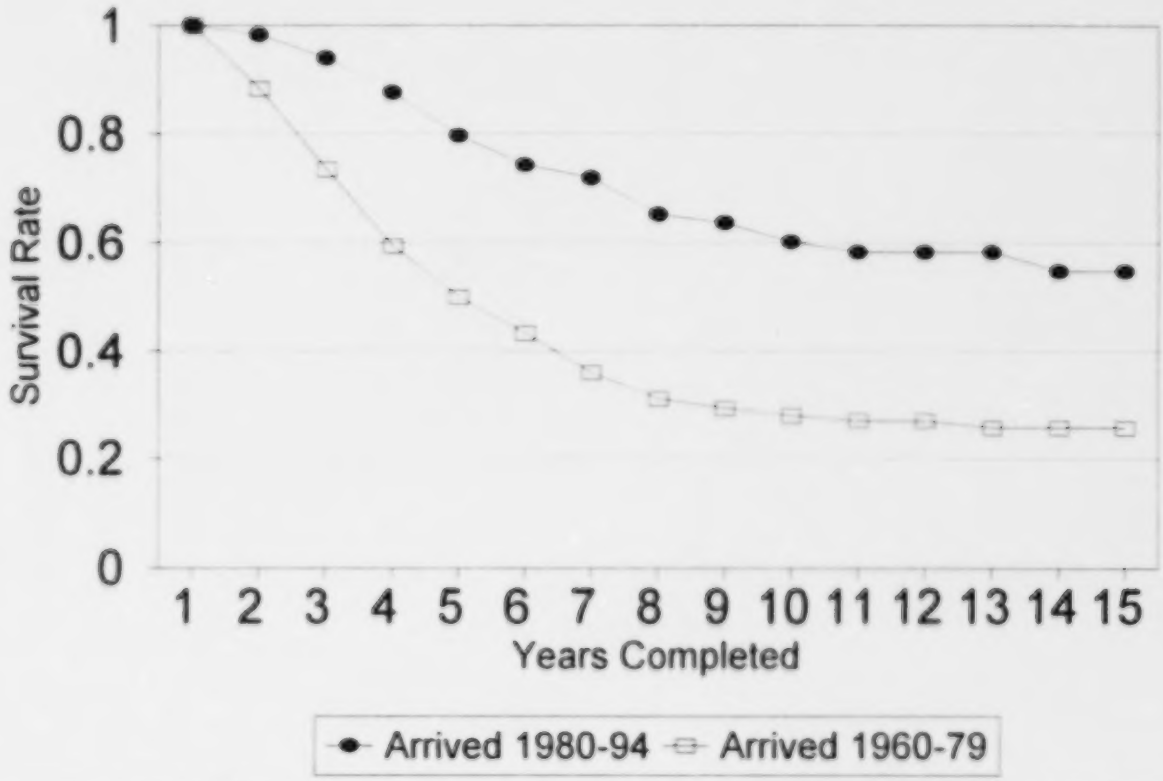

h:\faculty\survival\exhibit1.wb1

EXHIBIT 5: HAZARD RATES (% of survivors leaving prior to year X)

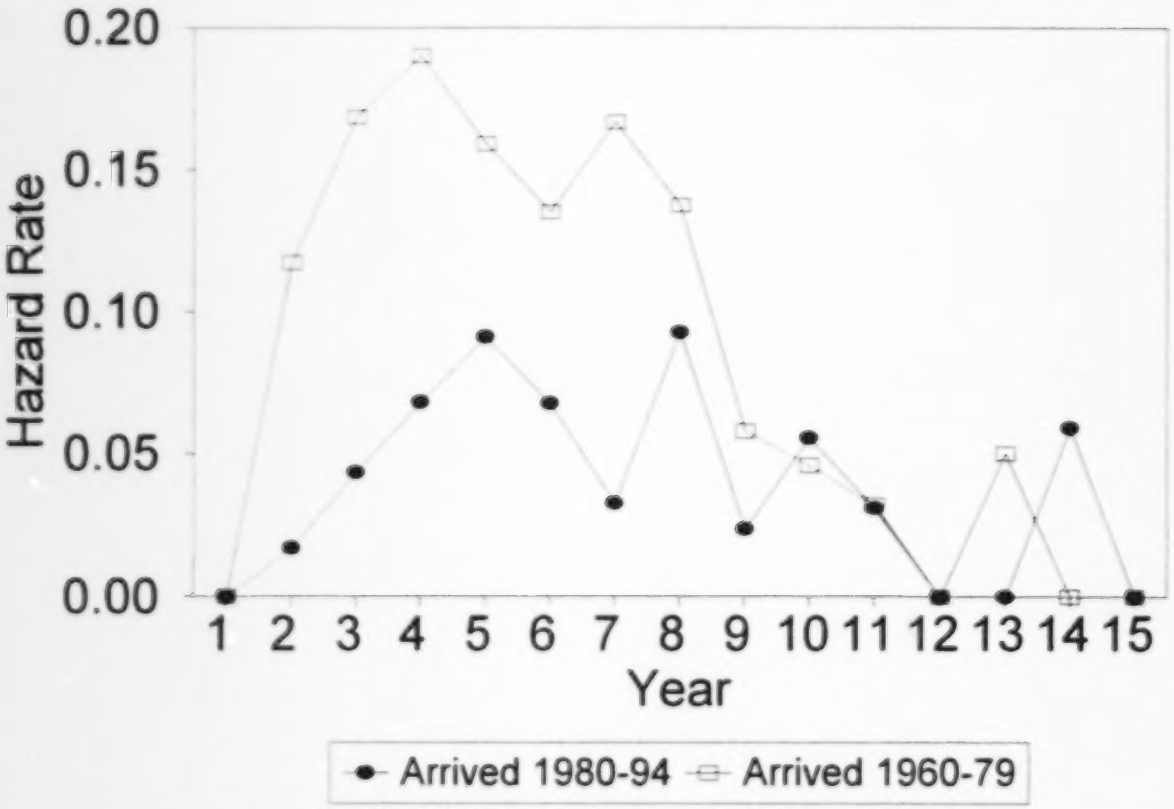

h:\faculty\survival\exhibit4.wb1

Exhibit 6: The Survival and Hazard Rates of Women Faculty Arriving 1980-94

Years Completed	Number of Faculty	Terminal Events	Censored	Propn Surviving	Cumul Propn Surv at End	Hazard Function
1	52	0	0	1.000	1.000	0.000
2	52	1	0	0.981	0.981	0.019
3	51	4	4	0.922	0.904	0.078
4	43	3	7	0.930	0.841	0.070
5	33	4	0	0.879	0.739	0.121
6	29	2	4	0.931	0.688	0.069
7	23	0	1	1.000	0.688	0.000
8	22	2	4	0.909	0.625	0.091
9	16	0	2	1.000	0.625	0.000
10	14	1	1	0.929	0.581	0.071
11	12	0	2	1.000	0.581	0.000
12	10	0	1	1.000	0.581	0.000
13	9	0	2	1.000	0.581	0.000
14	7	1	1	0.857	0.498	0.143
15	5	0	2	1.000	0.498	0.000

Exhibit 7: The Survival and Hazard Rates of Men Faculty Arriving 1980-94

Years Completed	Number of Faculty	Terminal Events	Censored	Propn Surviving	Cumul Propn Surv at End	Hazard Rate
1	65	0	0	1.000	1.000	0.000
2	65	1	0	0.985	0.985	0.015
3	64	1	3	0.984	0.969	0.016
4	60	4	1	0.933	0.905	0.067
5	55	4	6	0.927	0.839	0.073
6	45	3	4	0.933	0.783	0.067
7	38	2	4	0.947	0.742	0.053
8	32	3	3	0.906	0.672	0.094
9	26	1	3	0.962	0.646	0.038
10	22	1	1	0.955	0.617	0.045
11	20	1	4	0.950	0.586	0.050
12	15	0	2	1.000	0.586	0.000
13	13	0	3	1.000	0.586	0.000
14	10	0	2	1.000	0.586	0.000
15	8	0	3	1.000	0.586	0.000

EXHIBIT 8: Survival Rates of Male and Female Faculty Arriving After 1979

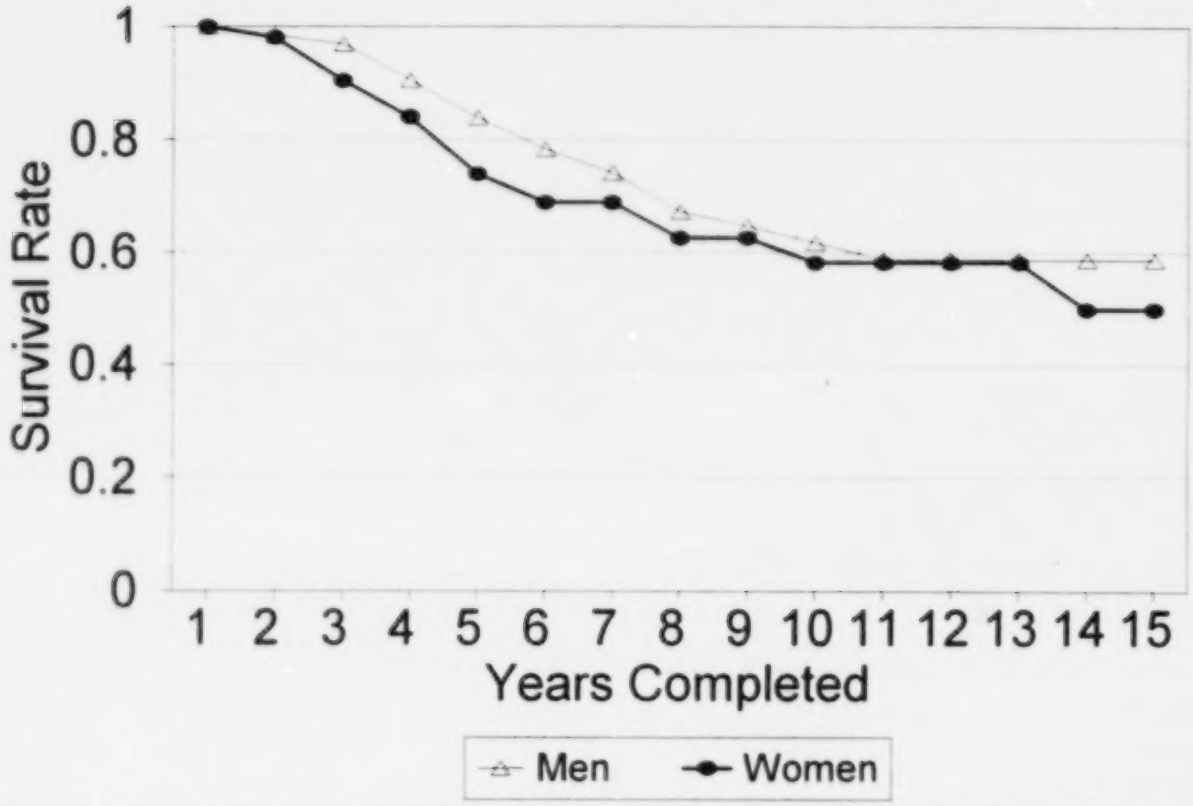

h:\faculty\survival\exhibit6.wb1

EXHIBIT 9. RESULTS FROM A COX PROPORTIONAL HAZARDS REGRESSION

Dependent Variable: Hazard Rate

Variable	Regression Coefficient	t-statistic
Female	.01	.04
Post1979***	-.72	-3.9
Tenured**	-.63	-2.0
HavePhD**	-.29	-2.0
Adjunct*	-.90	-1.8

*significant at the 10% level

**significant at the 5% level

***significant at the .1% level

338 observations, 118 censored, 220 exited

EXHIBIT 10: DO WE NEED TO STRATIFY?

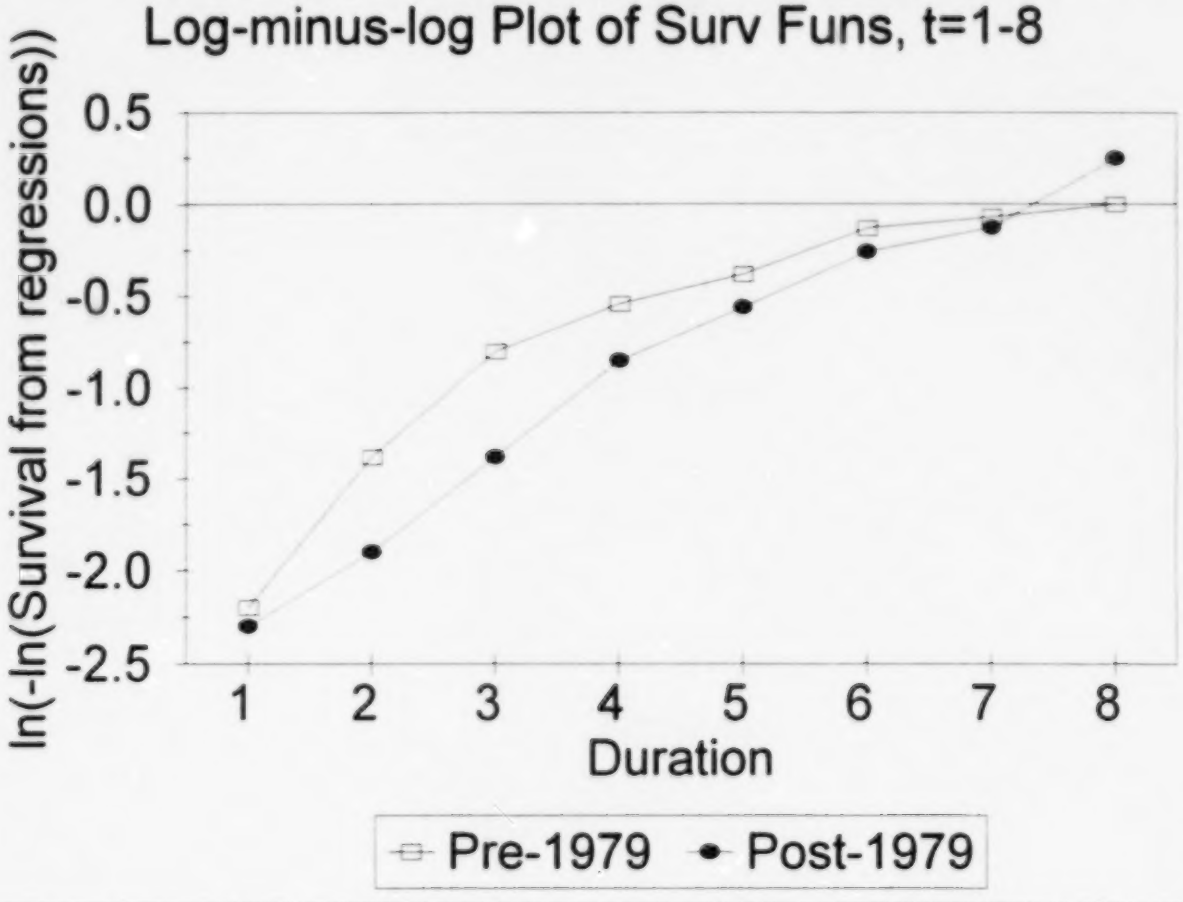

h:\faculty\survival\residmkt.wb1

EXHIBIT 11: GENERALIZED RESIDUALS

Unstratified Regression

Censored Observations Counted as Exits

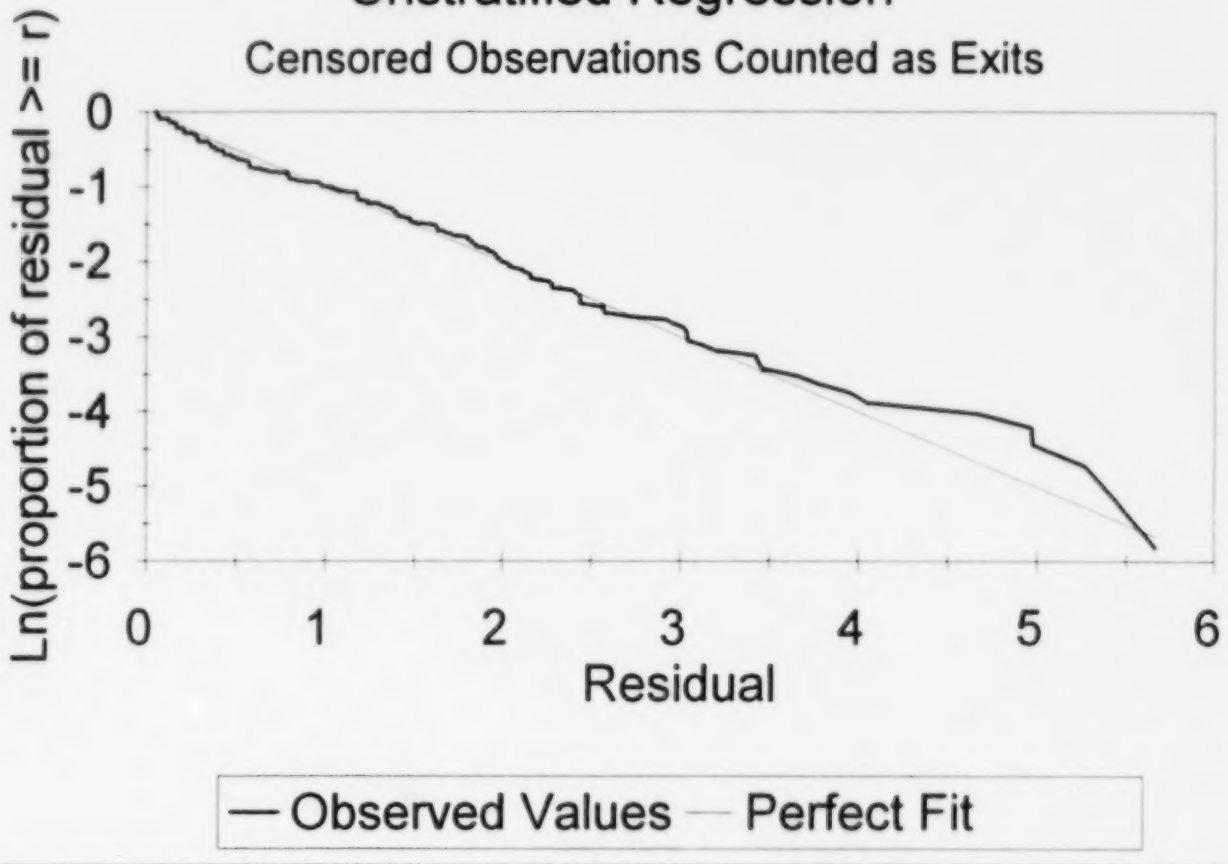

h:\faculty\survival\residmkt.wb1

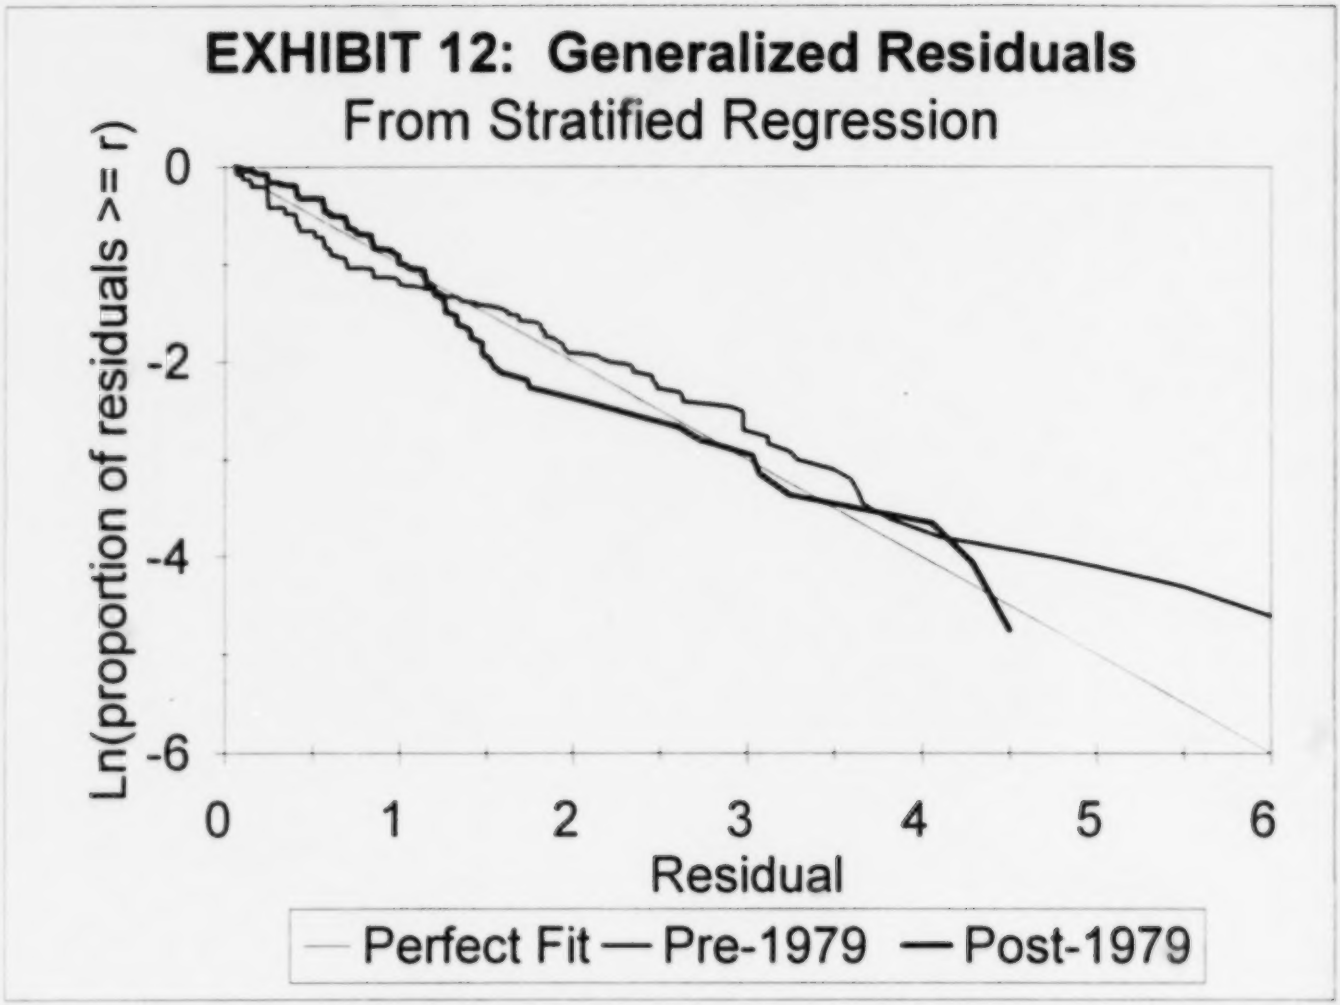

h:\faculty\survival\residmkt.wb1